

DUPLICATE

Nov 0 PCT/PTO

25 APR 2005

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

10/532609

(19) World Intellectual Property
Organization
International Bureau



(43) International Publication Date
13 May 2004 (13.05.2004)

PCT

(10) International Publication Number
WO 2004/040877 A1

(51) International Patent Classification⁷: H04L 29/06

(21) International Application Number:
PCT/GB2003/004374

(22) International Filing Date: 8 October 2003 (08.10.2003)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
0225356.5 31 October 2002 (31.10.2002) GB

(71) Applicant (for all designated States except US): **BRITISH TELECOMMUNICATIONS PUBLIC LIMITED COMPANY** [GB/GB]; BT Group Legal, Intellectual Property Department, PP C5A .BT Centre, 81 Newgate Street, London EC1A 7AJ (GB).

(72) Inventor; and

(75) Inventor/Applicant (for US only): **CLARK, Jonathan, Andrew** [GB/GB]; 10 SLEAFORD CLOSE, IPSWICH, Suffolk IP2 9PE (GB).

(74) Agent: **LIDBETTER, Timothy, Guy, Edwin**; BT GROUP LEGAL, INTELLECTUAL PROPERTY DEPARTMENT, HOLBORN CENTRE, 8th Floor, 120 Holborn, London EC1N 2TE (GB).

(81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.


(84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PT, RO, SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— with international search report

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: PARALLEL ACCESS TO DATA OVER A PACKET NETWORK



(57) Abstract: An end user application (15) generates a plurality of access requests for the same data to be delivered over a plurality of routes (21, 22, 23), each request conveying an indication of their common origin to the targeted webserver or other internet application (13). The targeted application (13) has means arranged to identify the indication of common origin, and therefore whether a plurality of addresses (21X, 22Y, 23Z) requesting multiple requests for the same data are associated with the same end user (15), and where this is the case splits the requested data, and streams different parts of the data to the different addresses requesting it. The end user, on receiving the requested data, assembles the data sent over the plurality of routes into a single stream for access by the user. Buffering may be necessary if traffic is slower over one path than it is over another.

WO 2004/040877 A1

PARALLEL ACCESS TO DATA OVER A PACKET NETWORK

This invention relates to the accession of data from a database, and in particular the accession of large files or streamed (non real time) video from remote
5 websites.

Currently, if a user having a high bit rate (2Mb/s) ADSL (asymmetric digital subscriber loop) connection attempts to "stream" a high bit-rate video computer file from a remote site, perhaps in another country, the bit rate received is often very much slower than the user's connection is capable of, perhaps only 100 - 400 kb/s,
10 due to congestion and contention for capacity with other users over shared international connections. At this rate video quality would be poor and virtually unacceptable. It is possible to alter the underlying network to provide say 2Mb/s throughput for the particular user, but this requires changes to the core Internet Protocol (IP) network. The concept of using multiple virtual connections and access
15 connections to support improved throughput is already established and implemented in the form of Inverse Multiplexing for ATM (IMA) (*ATM-Forum AF-PHY-0086.0010*): for internet access the end user would get a single IP address and the access server would reassemble the data into one stream for routing across the internet. Note the access server is the gateway between the connection orientated ATM access
20 network and the IP routed core. Therefore there would be no benefit if the bottleneck is in fact between the access server and the database being accessed.

Most access networks have the capability to allow an end user, if he so chooses, to gain access several times simultaneously, using the same or different Internet service providers. Each time the user logs on he gets another IP address. To
25 the core IP network these addresses all appear to be separate users. When it is heavily loaded, the core network divides its bandwidth up equally between the addresses contending for access, so a user who has logged on three times will get three times as much bandwidth as a user logged on only once. The mechanism to log on a plurality of times could be multiple asynchronous transfer mode (ATM)
30 permanent or switched virtual connections (PVCs or SVCs), or the PPPoE (Point to Point Protocol over Ethernet). However, if the internet application streaming the data to the end user receives requests for a particular stream from three separate IP address it will send the same data to all three addresses and so the additional

bandwidth will be of no practical use, as the data will simply be duplicated between the various links.

According to one aspect of the invention, there is provided a method of accessing data from an internet application over a distributed information network, wherein a user terminal generates a plurality of access requests for the same data to be delivered by the internet application over a plurality of routes, each request conveying an indication of their common origin to the targeted internet application, the internet application identifies whether a plurality of addresses making requests for the same data are associated with the same end user, and where this is the case splitting the requested data and streaming different parts of the data to the different addresses requesting it, and the user terminal receives the requested data over the plurality of routes and assembles it into a single stream.

According to another aspect of the invention, an internet application has means arranged to identify whether a plurality of addresses requesting multiple requests for the same data are associated with the same end user, and where this is the case splits the requested data, and streams different parts of the data to the different addresses requesting it. According to a complementary aspect, the end user application is provided with means for generating a plurality of access requests for the same data to be delivered over a plurality of routes, each request conveying an indication of their common origin to the targeted internet application, and means for receiving the requested data and to assemble the data sent over the plurality of routes into a single stream for access by the user.

Buffering may be necessary if traffic is slower over one path than it is over another.

In a preferred embodiment the internet application comprises means for identifying correlation codes associated with data requests, means for associating each such data request with any previous requests for the same data having the same correlation code, and means for splitting the requested data between the addresses associated with the data requests. The corresponding user terminal comprises means for generating a first access request having a correlation code indicative of its origin, means for determining whether the data rate of the data received in response to the first request meets a predetermined level, and means to

generate one or more further requests over different routes using the same correlation code.

The invention offers an improved quality of service and improved download speed. The invention requires the internet application and the user equipment to co-
5 operate such that the internet application can identify addresses of users making use of this invention, but requires no changes to the core internet or IP networks currently deployed, or their component equipment (Routers and Access Servers). The user may connect conventionally several times to the same Internet Service Provider, or may prefer to simultaneously connect to multiple Internet Service Providers – a
10 capability known to be possible with multiple virtual channels, either permanent or switched.

Therefore with a simple change to the broadband access network, the end user client software and internet application, the end user can receive non real time data at several times the rate of a user with only a single Internet connection. The
15 management and support systems of the access network may require modification to provide the ability to operate the user's broadband Internet connection as a plurality of virtual channels, which may each be connected to a different ISP.

The greater delay, and greater variation in that delay, may require larger initial buffers and may result in the video taking longer to start up and appear on the screen
20 after being requested. In one preferred arrangement a first stream is connected in the conventional way, others being added if the received bit-rate is not deemed adequate. Thus the initial delay would be minimised. Viewers who frequently switch between different TV channels would therefore be able to identify what is being shown on the channel without delay, but may have to wait for the quality to reach an optimum
25 level.

An embodiment of the invention will now be described, by way of example, with reference to the drawings in which

Figure 1 is a schematic illustration of a prior art conventional Single ISP Connection Service;

30 Figure 2 is a schematic illustration of the connection of three switched virtual connections (SVCs) to three Internet protocol (IP) Addresses;

Figure 3 is a schematic illustration of why streaming would not work over multiple conventional SVCs

Figure 4 is a schematic illustration of a system operating according to the invention, with correlated streaming over multiple paths.

Figure 5 is a flow chart illustrating the method of operation of one embodiment of the invention.

5 As shown in Figure 1, currently a single permanent virtual circuit (PVC) 10 is set up between the user terminal 1 of an ADSL user, and the access server 11. The access server 11 terminates the PVC 10 and the PPP (12) signalling encapsulated onto the PVC. It also gives the end user terminal 1 an IP address 1X to enable it to connect to the Internet 14 and send data to any other Internet application, for
10 example a broadcast webserver 13 (Figures 3 and 4).

 With a switched virtual circuit (SVC) ADSL multiplexer (DSLAM) 20 the end user terminal 15 can connect simultaneously to multiple access servers 21, 22, 23. Figure 2 shows the user logged on to three different Internet service providers. Again using PPP he is given a unique routable IP address 21X, 22Y, 23Z for each
15 connection. Although the increased number of connections marginally increases the contention level (the number of individual IP addresses attempting to access data), at busy times the individual user connected to three ISPs as in Figure 2 will still get almost three times as much data through as a user connected to a single ISP as in the example of Figure 1. (The actual increase in rate is $3n/(n+2)$ where n is the total
20 number of IP addresses accessing data. When traffic is busy, the value of " n " is large and so the expression converges on the value 3.0)

 If all three connections are made to the same destination server 13, the user 15 of the arrangement in Figure 2 will not gain in overall information rate, because most of the data will be duplicated. Figure 3 shows how, using the existing IP
25 network streaming protocols, the server 3 being accessed would consider the user 15 to be three different users because it receives requests from three different IP addresses 21X, 22Y, 23Z. It would then send the same data to the end user 15 over each of the three different routes, via the access servers 21, 22, 23, so there is clearly little benefit in setting up extra SVCs and trying to downstream the data over
30 them.

 Consider the scenario where the user 15 is trying to downstream video at 500kb/s from a broadcast webserver 13 in another continent. Each of the

connections 21, 22, 23 are heavily loaded and can each only supply 200kb/s. The user will obtain useful data only at the rate of the fastest of the three connections.

In the embodiment shown in Figures 4 and 5, an enhanced streaming protocol is provided, containing a correlation ID from the end user. The correlation ID is chosen such that it is unlikely to be duplicated by other users. It could be generated in a variety of ways, either randomly or based on the user's unique address..

As shown in Figure 5, the user terminal 15 logs on to several Internet Service Providers' access servers 21, 22, 23....., obtaining a different address X, Y, Z..... from each server (step 50). It would be possible to establish two or more connections to the same ISP. However, some ISPs have authentication systems designed to prevent simultaneous logons by the same user, in order to prevent fraudulent access. Such authentication systems would require reconfiguration to allow such simultaneous logons to take place.

The user terminal 15 first makes an initial streaming request 51, including the unique correlation code, over a first access server 21.

The broadcast webserver 13 checks this request against a store of previous requests (step 52) but fails to find any such requests with the same correlation code. Since this is the first request for this data that the user 1 has made, no such previous request has been recorded and the video stream is returned to the user 1 in the conventional manner (step 53).

The user terminal 15 now checks the data rate of the video stream against a predetermined value (step 54). If the data rate is too slow, the user terminal 1 transmits a similar request 55, using the same correlation code but using a different access server 22.

The user terminal may also start to show the video stream with the reduced quality dictated by the low bit rate, so that the user can see what is being received. Alternatively, the data may be buffered so that the stream can all be shown at full quality when the further stream or streams have been added. The reduction in quality of the first option is preferred when delay is undesirable, such as when a real-time signal is being transmitted, or if a user is sampling a number of feeds to see what is available.

The webserver 13 again checks this request 55 against the store of previous requests (step 52) but this time recognises that the requests 51, 55, despite coming from different IP addresses, 21X, 22Y are in fact from the same origin 1. The server 3 then apportions the data between the connections 21, 22 according to the rates they can each support (step 56). The information on attainable data rates can be obtained from, for example the TCP sliding window size in the current TCP/IP stack. The windows size adapts to the data throughput in the current Internet TCP session, , so it is a reasonably accurate representation of throughput. . A small data overhead is required in the transmitted data to identify the order in which the data is to be reassembled.

This process is repeated until the user terminal 1 determines (step 54) that the data rate is satisfactory (or all available addresses have been used), and then the received data is buffered and assembled in the correct order (step 57). Consequently, over several service providers 21, 22, 23 the user's effective data rate would be the sum of the service providers' throughputs, rather than just the fastest one on its own. So in this example three 200 kb/s connections would provide 600kb/s. Thus a 500kb/s TV stream could be supported by the three connections working together, where none of them could do so on its own. In order to avoid overloading of the network by users attempting vast numbers of parallel access attempts, the internet application may limit the number of connections available to any given user. However, in practice a user attempting to use more than a few connections would experience no greater benefit, as the bandwidth of his own access connection would become the limiting factor.

The invention may be used in conjunction with the invention described in the applicant's co-pending International application filed on the same date as the present application and claiming priority from United Kingdom patent application 0225359.9. That other application describes a method of improving the latency (delay) of a signal by transmitting it in its entirety over several parallel channels such that, for each packet sent to the destination, the first instance of that packet to arrive is assembled with the first instance of the other packets to arrive to form a single output stream. For example using six feeds (IP Addresses), a stream may be split into two to double the bandwidth according to the present invention, and then these two streams are then each duplicated three times to reduce delay.

CLAIMS

1. An internet application for providing data on receipt of requests from user terminals over a distributed information network, having means arranged to identify
5 whether a plurality of addresses making requests for the same data are associated with the same end user, and where this is the case splitting the requested data, and streaming different parts of the data to the different addresses requesting it.
2. An internet application according to claim 1, comprising means for identifying
10 correlation codes associated with data requests, means for associating each such data request with any previous requests for the same data having the same correlation code, and means for splitting the requested data between the addresses associated with the data requests
- 15 3. An internet application according to claim 1 or claim 2, comprising means for identifying the data rates available to each of the requesting addresses and apportioning the data between the addresses accordingly.
4. A user terminal for accessing data from an internet application over a
20 distributed information network, provided with means for generating a plurality of access requests for the same data to be delivered by the internet application over a plurality of routes, each request conveying an indication of their common origin to the targeted internet application, and means for receiving the requested data and to assemble the data sent over the plurality of routes into a single stream for access by
25 the user.
5. A user terminal according to claim 4, comprising means for generating a first access request having a correlation code indicative of its origin, means for determining whether the data rate of the data received in response to the first
30 request meets a predetermined level, and means to generate one or more further requests over different routes using the same correlation code.

6. A user terminal according to claim 4 or 5, comprising means for buffering the incoming data to allow its reassembly in a manner prescribed by the data content.

7. A method of accessing data from an internet application over a distributed
5 information network, wherein a user terminal generates a plurality of access requests for the same data to be delivered by the internet application over a plurality of routes, each request conveying an indication of their common origin to the targeted internet application, the internet application identifies whether a plurality of addresses making requests for the same data are associated with the same end user, and where this is
10 the case splitting the requested data and streaming different parts of the data to the different addresses requesting it, and the user terminal receives the requested data over the plurality of routes and assembles it into a single stream.

8. A method according to claim 7, wherein the user terminal generates an initial
15 access request with a correlation code indicative of its origin and the internet application stores the correlation code, and if the user terminal determines that the data received in response to the initial request does not meet a predetermined data rate, it transmits one or more further requests using the same correlation code, the internet application identifying such requests as being associated with the same end
20 user.

9. A method according to claim 7 or 8, wherein the internet application identifies the data rates available on the connection to each of the requesting addresses and apportions the data to be transmitted to each of the addresses
25 accordingly.

10. A method according to claim 7, 8 or 9, wherein the incoming data contains information to allow the user terminal to reassemble it, and the user terminal buffers the information to allow its reassembly accordingly.

1/5




Figure 1

2/5




Figure 2




Figure 3




Figure 4

